

**APPLICATION FOR
UNITED STATES PATENT
IN THE NAME OF**

Steven H. Gardner

OF

CONEXANT SYSTEMS, INC.

FOR

**Method and Apparatus for Medium Access Control in Powerline
Communication Network Systems**

DOCKET NO. CNXT-002-PAP - (CNXT REF. NO.: 00CXT0307-PCD)

Prepared by

Marty Jaquez and Allan Lee

Express Mail No. EK916662885US

Method and Apparatus for Medium Access Control in Powerline Communication Network Systems

CROSS-REFERENCE TO RELATED PROVISIONAL APPLICATION

5

This application claims the benefit of U.S. Provisional Application No. 60/210,148, filed June 07, 2000, entitled "Method and Apparatus for Medium Access Control in Powerline Communication Network Systems", hereby incorporated by reference herein.

10

BACKGROUND OF THE INVENTION

1. *Field of the Invention*

15

This invention relates to powerline communication networks, and more particularly to a method and apparatus for medium access control in powerline communication network systems.

2. *Description of Related Art*

20

The past few years have brought about tremendous changes in the modern home, and especially, in appliances and other equipment designed for home use. For example, advances in personal computing technologies have produced faster, more complex, more powerful, more user-friendly, and less expensive personal computers (PCs) than previous models. Consequently, PCs have proliferated and now find use in a record number of homes. Indeed, the number of multiple-PC homes (households with one or more PCs) is also growing rapidly. Over the next few years, the number of multiple-PC homes is expected to grow at a double-digit rate while the growth from single-PC homes is expected to remain flat. At the same time, the popularity and pervasiveness of the well-known Internet has produced a need for faster and less expensive home-based access.

25

30

As is well known, usage of the Internet has exploded during the past few years. More and more often the Internet is the preferred medium for information exchange, correspondence, research, entertainment, and a variety of other communication needs. Not surprisingly, home-

based Internet usage has increased rapidly in recent years. A larger number of homes require access to the Internet than ever before. The increase in home Internet usage has produced demands for higher access speeds and increased Internet availability. To meet these needs, advances have been made in cable modem, digital subscriber loop (DSL), broadband wireless, powerline local loop, and satellite technologies. All of these technologies (and others) are presently being used to facilitate home-based Internet access. Due to these technological advances and to the ever-increasing popularity of the Internet, predictions are that home-based Internet access will continue to explode during the next decade. For example, market projections for cable modem and DSL subscriptions alone show an imbedded base of approximately 35 million connected users by the year 2003.

In additions to recent technological advances in the personal computing and Internet access industries, advances have also been made with respect to appliances and other equipment intended for home use. For example, because an increasing number of people work from home, home office equipment (including telecommunication equipment) has become increasingly complex and sophisticated. Products have been developed to meet the needs of the so-called SOHO ("small office, home office") consumer. While these SOHO products tend to be less expensive than their corporate office product counterparts, they do not lack in terms of sophistication or computing/communication power. In addition to the increasing complexity of SOHO products, home appliances have also become increasingly complex and sophisticated. These so-called "smart" appliances often use imbedded microprocessors to control their functions. Exemplary smart appliances include microwaves, refrigerators, dishwashers, washing machines, dryers, ovens, *etc.* Similar advances have been made in home entertainment systems and equipment such as televisions (including set-top boxes), telephones, videocassette recorders (VCRs), stereos, *etc.* Most of these systems and devices include sophisticated control circuitry (typically implemented using microprocessors) for programming and controlling their functions. Finally, many other home use systems such as alarm systems, irrigation systems, *etc.*, have been developed with sophisticated control sub-components.

The advances described above in home appliance and equipment technologies have produced a need for similar advancements in home communication networking technology. As home appliances and entertainment products become increasingly more complex and sophisticated, the need has arisen for facilitating the interconnection and networking of the home appliances

and other products used in the home. Also, a need for distribution of entertainment media such as PC applications, audio streaming and voice telephony exists. One proposed home networking solution is commonly referred to as "Powerline Networking". Powerline networking refers to the concept of using existing residential AC power lines as a means for
5 networking all of the appliance and products used in the home. Although the existing AC power lines were originally intended for supplying AC power only, the Powerline Networking approach anticipates also using the power lines for communication networking purposes. One such proposed powerline networking approach is shown in the block diagram of FIGURE 1.

10 As shown in FIGURE 1, the powerline network 100 comprises a plurality of power line outlets 102 electrically coupled to one another via a plurality of power lines 104. As shown in FIGURE 1, a number of devices and appliances are coupled to the powerline network via interconnection with the plurality of outlets 102. For example, as shown in FIGURE 1, a
15 personal computer 106, laptop computer, 108, telephone 110, facsimile machine 112, and printer 114 are networked together via electrical connection with the power lines 104 through their respective and associated power outlets 102. In addition, "smart" appliances such as a refrigerator 115, washer dryer 116, microwave 118, and oven 120 are also networked together using the proposed powerline network 100. A smart television 122 is networked via electrical connection with its respective power outlet 102. Finally, as shown in FIGURE 1, the
20 powerline network can access an Internet Access Network 124 via connection through a modem 126 or other Internet access device.

With multiple power outlets 102 in almost every room of the modern home, the plurality of power lines 104 potentially comprise the most pervasive in-home communication network in
25 the world. The powerline network system is available anywhere power lines exist (and therefore, for all intents and purposes, it has worldwide availability). In addition, networking of home appliances and products is potentially very simple using powerline networking systems. Due to the potential ease of connectivity and installation, the powerline networking approach will likely be very attractive to the average consumer. However, powerline
30 networking systems presents a number of difficult technical challenges. In order for powerline networking systems to gain acceptance these challenges will need to be overcome.

To appreciate the technical challenges presented by powerline networking systems, it is helpful to first review some of the electrical characteristics unique to home powerline

networks. As is well known, home power lines were not originally designed for communicating data signals. The physical topology of the home power line wiring, the physical properties of the electrical cabling used to implement the power lines, the types of appliances typically connected to the power lines, and the behavioral characteristics of the current that travels on the power lines all combine to create technical obstacles to using power lines as a home communication network.

The power line wiring used within a house is typically electrically analogous to a network of transmission lines connected together in a large tree-like configuration. The power line wiring has differing terminating impedances at the end of each stub of the network. As a consequence, the transfer function of the power line transmission channel has substantial variations in gain and phase across the frequency band. Further, the transfer function between a first pair of power outlets is very likely to differ from that between a second pair of power outlets. The transmission channel tends to be fairly constant over time. Changes in the channel typically occur only when electrical devices are plugged into or removed from the power line (or occasionally when the devices are powered on/off). When used for networking devices in a powerline communications network, the frequencies used for communication typically are well above the 60-cycle AC power line frequency. Therefore, the desired communication signal spectrum is easily separated from the real power-bearing signal in a receiver connected to the powerline network.

Another important consideration in the power line environment is noise and interference. Many electrical devices create large amounts of noise on the power line. The powerline networking system must be capable of tolerating the noise and interference present on home power lines. Some of the home power line interference is frequency selective. Frequency selective interference causes interference only at specific frequencies (*i.e.*, only signals operating at specific frequencies are interfered with, all other signals experience no interference). However, in addition, some home power line interference is impulsive by nature. Although impulsive interference spans a broad range of frequencies, it occurs only in short time bursts. Some home power line interference is a hybrid of these two (frequency selective and impulsive). In addition to the different types of interference present on the home power lines, noise is neither uniform nor symmetrical across the power lines. For example, noise proximate a first device may cause the first device to be unable to receive data from a second, more distant device; however, the second device may be able to receive data from the

first. The second device may be able to receive information from the first because the noise at the receiver of the second device is attenuated as much as is the desired signal in this case. However, because the noise at the receiver of the first device is not as attenuated as is the desired signal (because the noise source is much closer to the first device than the second), the first device will be unable to receive information from the second.

Another consideration unique to powerline networking systems is that home power line wiring typically does not stop at the exterior wall of a house. Circuit breaker panels and electric meters (typically located outside the home) pass frequencies used for home networking. In typical residential areas, a local power transformer is used to regulate voltage for a fairly small number of homes (typically between 5 and 10 homes). These homes all experience relatively small amounts of attenuation between each other. The signal frequencies of interest to powerline networking systems do not tend to pass through the transformer. Due to these electrical characteristics, signals generated in a first home network can often be received in a second home network, and *vice versa*. In addition, unlike internal dedicated Ethernet or other data networks, power lines are accessible from power outlets outside of the home. This raises obvious security concerns because users typically do not want to share information with unauthorized users including their neighbors.

Signals that travel outside of the house tend to encounter greater attenuation than those that originate in the same house, and thus the percentage of outlets having house-to-house connectivity is much lower than the percentage for same house connectivity. The fact that transmissions at some outlets may not be receivable at other outlets is a significant difference between powerline networking systems and a wired LAN-type communication network such as the well-known Ethernet.

Despite these and other technical concerns, powerline communication network systems are presently being developed and proposed. For example, the HomePlug™ Powerline Alliance has proposed one such powerline communication network. The HomePlug™ Powerline Alliance is a non-profit industry association of high technology companies. The association was created to foster an open specification for home powerline networking products and services. Once an open specification is adopted, the association contemplates encouraging global acceptance of solutions and products that employ it.

A very important aspect of any home powerline networking system specification is the definition of a Medium Access Control ("MAC") communication protocol. The MAC protocol should be designed to allow devices to share the powerline network in a fair manner that provides performance in terms of delivered throughput, latency, and acceptable errors.

5 The MAC should facilitate powerline networking system performance suitable for a number of applications including file transfer, voice, networked gaming, and streaming audio/video. The MAC protocol should be designed specifically to address the technical challenges posed by powerline networking systems. For example, the MAC communication protocol for powerline networking systems should ease compatibility of upgraded devices and protocols.

10 That is, the MAC should ease the efforts associated with installing and operating upgraded (*i.e.*, newer version) powerline networking system protocols and devices. Heretofore, it has been very difficult (if not impossible) to operate newer version devices on powerline networking systems operating older version devices. Therefore, there is a need for a MAC protocol that eases the task of upgrading powerline networking system protocols and devices.

15 In addition, the prior art attempts do not provide a mechanism for establishing "circuit-like" connections wherein devices are connected together via a "virtual circuit." Most prior art networking systems use a "contention-based" access approach. In this mode of access, when a device has data to send, it first determines whether a channel is being used by another device, and if it is not, the device begins data transmission. Other devices refrain from transmitting on the channel until the first device terminates its transmission. When more than one device requires transmission, a "collision" occurs and the devices re-transmit their data because procedures defined in the prior art MAC protocol require data re-transmission. As traffic on the channel increases, so too do the number of collisions, resulting in data transmission delays.

20 The exact period of the delay is not fixed, but varies depending upon traffic characteristics. These delays may be acceptable for some applications (such as file transfers), but not others (such as streaming audio/video). The prior art solutions include use of priority schemes whereby "real-time" applications are assigned higher priorities than non-real-time applications. Disadvantageously, this approach is not ideal when there are multiple high

25 priority users as they must still contend with one another and thus still encounter probabilistic delays. This prior art approach also disadvantageously allows a low priority device to monopolize use of a channel (once it gains use of the channel) even when there are higher priority devices waiting to use the channel.

30

Therefore, there is a need for a MAC protocol that overcomes the disadvantages associated with the prior art solutions. A need exists for a MAC communication protocol that permits the reservation of transmission time between devices requiring "circuit-like" connections, and especially in environments wherein no central controller is used (such as those contemplated for use in powerline networking systems). A need exists for a MAC that facilitates the establishment of "virtual circuits" between devices in a powerline networking system. The MAC should tightly control transmission delays in a powerline networking system.

Further, any MAC designed for use in powerline networking systems should provide for the management and distribution of encryption keys, even in network environments having no central controller. For example, the powerline networking systems currently being proposed and developed (such as the powerline network 100 of FIGURE 1) do not contemplate use of central network controllers. One important aspect of MAC protocols designed for use in powerline networking systems is the control and distribution of encryption keys, especially in a controller-less network environment. In devices having user input/output (I/O) capabilities, encryption keys (or passwords that can be converted into encryption keys using a "hashing algorithm" or similar means) can be manually entered from the device. However, many devices may not have user I/O capability therefore making manual key entry impossible. Therefore, a need exists for a MAC communication protocol for powerline networking systems that facilitates the control and distribution of encryption keys, including the management of encryption keys for devices not having user I/O capability.

Finally, any MAC designed for use in powerline networking systems should provide for the unique assignment of logical network identifiers ("LNI"). Although all of the devices in a powerline networking system share the same physical medium (*i.e.*, the home power lines such as the power lines 104 of FIGURE 1), it is desirable to be able to separate the devices into logical networks ("LN") wherein only those devices belonging to the same logical network are allowed to share data. Using encryption schemes, data can be shared between devices that are members of a given LN, but is protected from devices that are not members of the given LN. It is very convenient and efficient to permit a device in an LN to determine which LN it belongs to. The device can use this information to determine whether it should attempt to receive a given data packet. For example, this can be accomplished by including the LNI in each data packet transmission. Alternatively, this can be accomplished by including a management message in which the transmitting device indicates its LNI. Under

In powerline networking systems wherein no central controller exists, there is no convenient means for ensuring that LNIs are not accidentally re-used. For example, it is not likely that home owners would feel comfortable asking their neighbors which LNI they have selected for their LN. Therefore, a need exists for a MAC protocol for a powerline networking system that facilitates the assignment of unique Logical Network Identifiers that are not accidentally re-used by the system.

The present invention provides such a method and apparatus for Medium Access Control in powerline communication network systems.

SUMMARY OF THE INVENTION

5 The present invention is a novel method and apparatus for Medium Access Control (MAC) of a physical medium in a powerline networking system. Several novel aspects of the present MAC method and apparatus are described including: use of a blanking interval to ease backward compatibility; use of beacons to propagate interval timing information throughout the network; controller-less reservation-based access used to establish and maintain "virtual circuit" connections between two selected devices; methods for generating and assigning Logical Network Identifiers (LNIs) in a powerline networking system; and a method for
10 generating, managing and distributing encryption keys, especially for devices having no user input/output capability.

15 The inventive MAC method and apparatus eases backward compatibility of the powerline networking system by providing a blanking interval method. In accordance with this inventive approach, blanking intervals define when previous version devices are prevented from transmitting data in the powerline network. The method described herein enables newer version devices to specify to older version devices the time periods that are reserved for use only by the newer version devices. In accordance with this approach, newer version devices first determine which newer version device controls the blanking interval timing. The selected
20 controlling device then transmits a "medium blanking payload" message containing information that specifies when the blanking interval occurs. The blanking interval is periodic. However, its period and duration can be varied by the controlling device. Thus, the blanking interval can adapt to network traffic conditions and requirements.

25 A beacon method is described that is closely tied to the inventive blanking interval method and apparatus. In accordance with the present invention, a method of using beacons is provided that is used to propagate blanking interval timing information throughout the powerline networking system. Using a "lifetime" mechanism, devices are able to determine whether and when blanking interval timing information becomes obsolete. The beacons are
30 also used by the devices to specify certain capabilities associated with the devices. The beacons are used to broadcast specific network capabilities associated with each device transmitting the beacon.

The present MAC protocol method and apparatus also provides a method for establishing "virtual circuit" connections between devices, wherein the virtual circuit connections have very tightly controlled throughput, delay, and latency characteristics. The inventive MAC protocol for powerline networking systems supports controller-less reservation based access modes that permit the creation of virtual circuit connections that provide periodic, low latency, constant bit-rate service between an originating device and a destination device. The virtual circuits are created by establishing a periodic time slot that is reserved for use only by a specific transmitter. All other devices in the network must be aware of the reservation and must avoid transmitting during the reserved time periods. The inventive method can be used in networks that have no centralized control mechanism.

In accordance with the present invention, devices can be logically separated into logical networks, wherein only devices belonging to a logical network are allowed to freely exchange information. The present invention includes a method for generating and assigning logical network identifiers (LNIs) that uniquely identify the logical networks in the system. The devices use the LNI information when determining whether to access information transmitted in a given data packet. A novel method of assigning LNIs to a logical network based upon user-selected password information is described. In one exemplary embodiment described herein, a character-based network name of variable length is used. The network name may be related to the user's street address. The name is compressed to a unique LNI using some form of a hash function. The method is particularly applicable to powerline networking systems having no central control mechanism.

Finally, a method of encryption key management and distribution is described, especially for devices having no user input/output (I/O) capability. For devices lacking user I/O capability manufacturers provide a hard-wired encryption key or default key that is not meant to be changed by users. However, the default key can be changed to be any possible value through utilization of a powerline-networked device that has user I/O capability. During an installation process, a system user inputs the default key (or a password from which the default key has been derived) that is associated with the new device (*i.e.*, device that lacks user I/O capability) into a powerline-networked device that has user I/O capability. An application program that is suited for this task aids the device during the installation process. The device that has I/O capability utilizes the default key to encrypt the current logical network key. The inventive MAC protocol then transmits the encrypted key to the new device (*i.e.*, device

without I/O capability). The logical network key is thus securely passed to the new device and all other members of the logical network thereafter exchange encrypted data with the device. If the device loses the logical network key, or if the key changes, another device re-transmits the key using the same MAC management message originally used to provide the

5

key.

2009-11-14 14:00:00

BRIEF DESCRIPTION OF THE DRAWINGS

FIGURE 1 shows an exemplary home powerline network system including a plurality of power line outlets electrically coupled to one another via a plurality of power lines.

5

FIGURE 2 is a timing diagram showing the timing of the reservation establishment procedure in accordance with present invention.

10

FIGURE 3a is a message flow diagram showing the three-way handshake used to renew or terminate reservations in accordance with the reservation renewal process of the present invention.

15

FIGURE 3b is a timing diagram showing the timing used by the reservation renewal packets in accordance with the present invention.

FIGURE 4 shows the Cyclic Redundancy Check (CRC) generator used to generate Logical Network Identifiers (LNI) in accordance with one embodiment of the present invention.

20

Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION OF THE INVENTION

Throughout this description, the preferred embodiment and examples shown should be considered as exemplars, rather than as limitations on the present invention.

5

The present invention is a method and apparatus for Medium Access Control (MAC) in powerline communication network systems. As described above, in home powerline networking systems, a single physical medium (the home power lines) is shared by a number of different devices (also referred to herein as "clients"). It may be desirable to allow some clients or devices to freely exchange data between them (for example, clients in a common household). In contrast, it may be desirable to prohibit certain clients from exchanging data with other clients (for example, it may be desirable to prevent a PC in a first house from exchanging data with a PC in a second house). The MAC functions to ensure that the physical medium (the home power lines) is shared in a fair, consistent, and efficient manner (*i.e.*, at a high performance level). The performance of the powerline networking system must be sufficient to accommodate a wide variety of clients and applications, including, but not limited to, file transfer, voice, networked games, and streaming audio/video applications.

10

15

20

25

The efficiency at which communication systems use the shared communication medium is a very important performance criterion in any communication system, especially in communication systems that have a physical communication medium shared by a plurality of differing devices or clients. Because powerline networking systems are, by definition, shared-medium communication networks, access and transmission by clients in the network must be controlled. The MAC protocol is used for this purpose to control client access to the physical communication medium. The MAC determines when clients are allowed to transmit on the physical medium and when they are not allowed to transmit. In addition, if contentions are permitted, the MAC controls the contention process and resolves any collisions that may occur.

30

With the various types of service applications from different devices, the powerline networking system MAC should adhere to a communication protocol that minimizes contention between the devices and that allows service applications to be tailored to the delay and bandwidth requirements of each application and service. The MAC is part of a "layered" data transport protocol wherein the lowest data transport layer is the physical signaling layer.

The physical transport layer is used to interface the higher communication protocol layers with the shared physical medium. The MAC method and apparatus of the present invention is intended for use with a powerline networking physical data transport layer.

5 Many different types of physical layer signaling can be used with the present inventive MAC protocol layer. One exemplary approach uses a modulation referred to as Orthogonal Frequency Division Multiplexing or OFDM. This modulation achieves high bit-rates by dividing the frequency spectrum into a large number of very narrow frequency bands. Low bit-rate transmissions are transmitted using a separate carrier frequency in each of the narrow frequency bands. The appeal of this approach is that when a channel has frequency selective impairments, the protocol can be designed to determine if there are frequencies that are severely impacted on a link between two clients. If so, the clients can agree not to use those frequencies. Standard modulation schemes using adaptive equalization can also be used with the present inventive MAC without departing from the scope or the spirit of the present invention. Although the present inventive MAC protocol is intended to operate with a physical layer that uses OFDM as the underlying modulation scheme, many key aspects of the protocol are contemplated for use with other types of physical layers.

20 One exemplary embodiment of the present MAC method and apparatus uses a Carrier Sense Multiple Access (CSMA) protocol with modifications to support special requirements for applications having low latency requirements. The protocol supports both contention-based and reservation-based access schemes. Reservation-based access schemes operate in either a controller-less mode or in a mode wherein a network controller is used.

25 One exemplary embodiment of the present inventive MAC method and apparatus is described in detail in the attached Appendix A (entitled "Procedures for Medium Access Control in the PL Network System" and hereinafter referred to as the "MAC specification"). The attached MAC specification describes procedures for medium access control (MAC) including procedures for: contention-based access, beacons and medium blanking intervals, controller-less reservation based access, payload formats including definitions of pre-defined payload formats, encryption, requesting and transmitting test messages, and controller-based reservation access. Exemplary pre-defined packet formats are also described in the attached Appendix A. For example, contention-based access packets, reservation-based access packets,

acknowledgment packets, preambles, FEC coding, tone masks and tone maps are described in detail.

5 Several aspects of the inventive MAC method and apparatus as described in the attached
MAC specification of Appendix A provide advantages over the prior art solutions. Each of
these aspects is described in detail below. Although an exemplary embodiment of the present
MAC method and apparatus is described, those skilled in the art shall recognize that
modifications can be made to the described embodiments (and in the attached MAC
specification) without departing from the scope or the spirit of the present invention.
10 Therefore, the present invention is not limited by the examples given below, but only by the
scope of the claims. The concept of providing a blanking interval and a means for providing
beacons is now described.

Blanking Interval – Backward Compatibility of the Powerline Networking System

15 As described above in the background of the invention, the powerline networks differ from
closed medium communication networks (such as the well-known Ethernet data
communication network typically used in an office environment) and from existing telephone
networks. Backward compatibility of devices is a desirable, but not a critical feature of these
existing communication systems. If sufficient performance improvement can be achieved
20 simply by upgrading all of the devices in the network, customers usually will be persuaded to
replace all of the network interface devices in the network. Unfortunately, an analogous
upgrade option is unavailable in the proposed powerline networking systems. Because the
physical medium (home power lines) is typically shared between neighboring homes, an
inhabitant of a first home typically cannot simply replace or upgrade the hardware in a second
25 neighboring home in order to upgrade the powerline networking system. Consequently,
heretofore, network and device upgrades have had to be fully interoperable with prior
generation hardware (*i.e.*, devices designed for use with previous versions of the powerline
networking system). This typically has been a very restrictive and limiting system
requirement. Thus, backward compatibility is an important consideration in the development
30 of any MAC for use in powerline networking systems.

The inventive MAC method and apparatus addresses this backward compatibility issue to
facilitate fully interoperable upgrades with previous version devices and systems. As
described in more detail below, the present invention provides a blanking interval during

which devices compatible with newer version protocols (*e.g.*, devices using version 1.1 of the MAC protocol) can clear out (or nullify) devices compatible with older version protocols (*e.g.*, devices using version 1.0). The use of blanking intervals (coupled closely with the use of "beacons" described in more detail below) greatly eases backward compatibility of an upgraded protocol when new protocols are designed.

The blanking interval operates to allow later version devices to specify to earlier version devices selected time periods (blanking intervals) during which only the later version devices are allowed to communicate. The blanking intervals are reserved for communication among the later version devices only. In the embodiment described in the attached MAC specification, the blanking structure comprises repeated sequences of times during which earlier version devices are restricted from contention-based access ("the blanking period") and times when they are allowed contention-based access (the "v1.0 period"). Reservation based access by earlier version devices is allowed during the blanking period, but the reservation establishment must be initiated during the v1.0 period. Any protocol for channel access can be used during the blanking intervals as long as the protocol used confines its transmissions to the blanking interval. The inventive blanking interval technique advantageously allows protocols that are presently completely undefined to be backward compatible with existing protocols. One embodiment of the present blanking interval method is described in detail in the MAC specification of Appendix A in sections 2.2 (pages 8-10) and 3.1.2.1.1.7 (pages 29-30).

In the embodiment of the blanking interval method described in the attached MAC specification, newer version devices (referred to in the MAC specification as "non-v1.0 devices") first determine between themselves which newer version device will control the blanking interval. This newer version device specifies the blanking structure by transmitting a message referred to as a "medium blanking payload." The medium blanking payload message contains information that specifies when the blanking interval occurs. The format and fields used by one embodiment of the medium blanking payload is described at section 3.1.2.1.1.7 (pages 29-30) of the attached MAC specification. As described therein, although the blanking interval is periodic, its period and duration are parameters that can be controlled and adjusted by the controlling device. Thus the blanking interval can be adapted to meet traffic demands.

For example, the device controlling the blanking interval can monitor traffic during both the blanking and unblanked intervals. The controlling device can adjust the duration of the blanking interval (for example, by modifying the duration of a blanking time field of the medium blanking payload) to provide more or less bandwidth to previous version or later version devices, as dictated by the monitored traffic characteristics. Under typical operating environments, the blanking interval period may be set at 100 milliseconds, and the blanking interval duration may be set at 50 milliseconds. Although these exemplary values for the blanking interval will vary substantially, values of this scale result in negligible increased delay for most PC-based applications.

In the embodiment described in the attached MAC specification, the controlling non-v1.0 device must re-transmit the medium blanking payload message at least once every five seconds. The device transmits a "ROBO" mode broadcast packet that contains the medium blanking payload message providing a network timing reference and the timing of the blanking period as described above. In the exemplary embodiment, a special contention resolution slot is provided for use by non-v1.0 devices at the conclusion of each blanking period. The contention resolution slot is advantageously used to ensure that non-v1.0 devices transmit the blanking information without collisions. As described in the attached MAC specification, in the exemplary embodiment, the timing associated with the transmission of the medium blanking messages is tied to the timing of the blanking intervals to ensure that the transmit time for the blanking messages is reserved. This advantageously prevents the transmission of another device from colliding with the medium blanking message, and thus greatly increases the probability that the medium blanking message is received.

The blanking interval method and apparatus of the present MAC protocol provides yet another advantage when used in a controller-less powerline networking system, and when used with devices requiring relatively low delays. In situations where v1.0 devices establish controller-less reservations, the controller-less reservations are allowed to continue through the duration of the blanking intervals. In accordance with this approach, the non-v1.0 devices (*i.e.*, the devices operable with later versions of the protocol) must respect the reservations established by these v1.0 devices. Advantageously, this approach ensures that v1.0 devices requiring very low delays (*e.g.*, v1.0 devices providing streaming audio/video information) will experience low delays even when blanking intervals are occurring.

Random Back-off Method for Preventing Collisions at the End of the Blanking Interval

5 The present invention advantageously reduces the likelihood of collisions of v1.0 device transmissions (*i.e.*, transmissions from devices designed to operate with earlier versions of the MAC protocol) occurring at the end of the blanking intervals. In powerline networking systems, data packets are typically presented to the device network transmitters by processes operating at higher layers of the protocol. The processes that cause the various devices to transmit (such as file transfer processes, web browsing commands, e-mail processes, *etc.*) typically have no relationship to the physical and MAC layer protocols of the powerline networks. Therefore, the times at which data packets are presented to the device network transmitters by the higher layer processes are independent of the timing of the blanking intervals. If a packet arrives at a v1.0 network transmitter from the higher protocol layers during the blanking interval, the packet is queued for transmission until the blanking interval ends. The longer the duration of the blanking interval, the greater the likelihood that one or more transmitters will have packets queued when the blanking interval ends. To reduce packet collisions from occurring when the blanking interval ends, the protocol does not simply allow all of the network transmitters with queued packets to transmit at the end of the blanking interval. Instead, the MAC protocol uses a "random back-off method" for the transmission of queued packets.

10 20 In accordance with the random back-off technique, each transmitter having a queued packet selects a random number "M" for use in transmitting queued packets. In one embodiment, M is defined to be an integer between 1 and some maximum value "max_slots." When the blanking interval ends, the transmitter waits for a period of time equal to M time slots. In one embodiment, the time slots are approximately 30 microseconds in duration. Alternative duration time slots can be used without departing from the scope or spirit of the present invention. If the channel is available after waiting for M time slots, the transmitter begins transmitting its queued packets. If the channel is unavailable, the MAC protocol enters the "BACKOFF state" to transmit the queued packets. The BACKOFF state is defined in detail in the attached MAC specification at section 2.1.5 (pages 7 and 8), and therefore is not defined in more detail herein.

25 30 In one embodiment of the present MAC protocol method and apparatus, the value max_slots is contained in the medium blanking message. In this embodiment, the value of max_slots is

determined by the controlling non-v1.0 device and is based upon the duration of the blanking interval and the amount of v1.0 traffic anticipated in the powerline networking system. By allowing the max_slots value to be programmable and variable, and by tying the value of max_slots to the duration of the blanking interval and to the traffic volume, the present invention advantageously facilitates adjustment of the contention period based upon the traffic characteristics of the network. In an alternative embodiment where this flexibility is not required, max_slots comprises a pre-determined fixed value.

Because some clients or devices in the powerline networking system may be unable to successfully receive packets that contain the medium blanking payload messages, the inventive MAC protocol includes the capability of propagating the blanking information using "beacon" messages. Some nodes in a powerline networking system may be unable to receive medium blanking messages because channel conditions between these nodes and the controlling device are severely degraded. Therefore, the present invention provides beacons that allow devices on the network to propagate blanking information to nodes that are unable to receive medium blanking messages directly from the controlling device. As described in more detail below (and in section 2.2, at pages 9-10 of the attached MAC specification), the blanking messages are propagated through the powerline network using a "relay" technique wherein the blanking information is transmitted by other nodes in the network. The inventive technique of using beacon payload messages to propagate the blanking information throughout the powerline network is now described.

Beacon Messages – Method and Apparatus for Propagating Blanking Information throughout the Powerline Network

Although the concept of using beacon messages to propagate blanking interval information is closely related to the blanking interval method and apparatus described above, the beacon messages also serve additional purposes in the present MAC protocol. The beacons are management messages that are periodically transmitted by each v1.0 node on the network. The beacon messages are used by the powerline networking system to propagate information to all nodes on the network, including nodes that are unable to receive (due to channel degradation, interference, *etc.*) transmissions from every other node on the network.

As defined in more detail in the attached MAC specification, in one embodiment of the present invention, each device transmits a respective and associated beacon payload message

on a nominally periodic basis. The beacon payload message preferably contains a number of fields that serve various MAC protocol purposes. One exemplary embodiment of such a beacon payload message (and associated fields) is described in the attached MAC specification in section 3.1.2.1.1.1, at pages 23-24. Those skilled in the art shall recognize that alternative beacon messages (and alternative beacon message fields) can be used without departing from the spirit or the scope of the present invention.

As described in section 2.2 of the MAC specification, beacon payload messages are preferably transmitted in broadcast data packets. Each client (or device) in the powerline network transmits beacon packets at a nominal five-second rate to indicate its presence in the network, and to propagate system timing information throughout the network. If a device node is able to receive medium blanking messages from a non-v1.0 device, it simply uses the blanking information obtained from that device to assemble its beacon messages. However, if a device node is unable to receive medium blanking messages from the controlling non-v1.0 device, it must assemble its beacon message based upon information obtained from another source. In these situations, and in accordance with the present invention, the device node assembles its beacon messages based upon information received from the device beacon message containing the smallest "lifetime field" value.

In accordance with the present invention, the medium blanking messages include lifetime fields (also referred to as "logical distance" fields in the attached MAC specification). As described in detail in the MAC specification, the device nodes use the lifetime field values to determine which beacon message to use when assembling their beacon messages. When a device node receives a medium blanking payload from a non-v1.0 device, it sets the lifetime field of its beacon payload message to zero. It also sets beacon fields for the duration of the interval blanking time, duration of the v1.0 time, and max_slots value equal to the values contained in the medium blanking payload message received from the non-v1.0 device. It also computes other system timing values as described in more detail in the MAC specification.

However, if the device node cannot receive the medium blanking messages from the controlling non-v1.0 device (due to channel degradation, interference or other factors), it sets the lifetime field to a non-zero value. In the exemplary embodiment described in the attached MAC specification, the device node prepares the contents of its beacon payload message using

information received from another device beacon only if it has not received a medium blanking payload message in the last five seconds. It shall be appreciated by those skilled in the art that alternative time periods can be used (*i.e.*, the device could wait for a longer or shorter time period) without departing from the spirit or scope of the present invention. The device node prepares the contents of its beacon payload message using the beacon received in the last recent five seconds having the lowest logical distance field (or lowest "lifetime field" value). If there are multiple received beacons having the same lowest lifetime field values, then the most recently received beacon is selected by the device node.

10 In accordance with the present beacon method and apparatus, the device node sets the lifetime field of its own beacon payload message to a value that is equal to one more than the lifetime field of the selected beacon (the beacon payload message from which the device node last obtained its blanking information). As described above, in the exemplary embodiment, the source used for the blanking information must be a message that was received within the past 15 five seconds. The lifetime field is also used by the present invention as a means for indicating when blanking interval information on the network has become obsolete.

For example, if the blanking interval ceases because the non-v1.0 device was removed from the network (or due to some other cause), the devices that were transmitting beacons with lifetime fields set to zero will no longer receive a medium blanking message. They will have to obtain their blanking interval information from another beacon and will thus set their lifetime fields to a value that is greater than or equal to one. In the next five-second interval, no device node will receive a beacon having a lifetime field value less than one, so the minimum lifetime value will increase to two. Every five seconds, the lifetime field value will increment until it reaches some pre-defined maximum value. In one embodiment of the present invention, the pre-defined maximum value for the lifetime field is seven. Alternative maximum values can be used without departing from the scope or spirit of the present invention. When the lifetime field value reaches its pre-defined maximum, the device nodes assume that no blanking intervals exist and that the v1.0 devices can access the channel at any time (*i.e.*, blanking intervals are not presently in use).

The beacons not only specify the period and duration of the blanking interval, but the exact timing as to when the blanking interval begins. Thus, the beacons contain fields that provide an absolute time at which the beacon messages are transmitted. The beacons also contain a

field that provides the time at which the next blanking interval will begin. This mechanism allows for the propagation of absolute system timing. As described in more detail in the attached MAC specification, the beacons also allow each device to specify certain capabilities (or limitations) associated with the device. For example, a powerline networking system may require that all device nodes accessing the system use BPSK and QPSK modulation. Further, the system may require that all of its devices allow 8-PSK and 16-QAM as options. The beacons described herein can be used to inform the system of the capabilities and limitations of associated devices. Using the information contained in the beacons, more capable nodes can be allowed to negotiate more efficient transmission modes when communicating on a point-to-point basis. In addition, the beacon method and apparatus can be used to negotiate alternative encryption algorithms.

Controller-less Reservation Based Access Method and Apparatus

As noted above in the background of the invention, the present MAC protocol method and apparatus includes a means for providing "virtual circuits" between devices connected to the powerline networking system. Most data communication networks (such as the well known Ethernet network) use "contention based access" techniques to control access to the network. In this mode of access, when a first client has data to transmit, it first checks the channel to see if it is presently being used, and if it is not being used, the first client initiates its transmissions. Other clients refrain from transmitting until this first client has finished transmitting its data. Occasionally, two or more clients may decide to transmit at the same time. In this case a "collision" will occur and typically neither client is successful in transmitting its data. Various schemes exist to resolve this situation. In most cases, the clients must eventually re-transmit their associated data until the data is successfully transmitted over the network. In order to reduce the likelihood of collisions occurring during the re-transmissions, many prior art systems use techniques to randomize the start time of the re-transmissions.

Disadvantageously, in the prior art network systems that use this type of access scheme, transmissions often do not occur immediately at the time when data is available for transmission. Indeed, when data traffic is heavy, it may take quite a while for a client to gain access to the channel. When access is finally obtained, collisions may further extend the time required to get the data through to its intended recipient. The exact delay encountered is not a fixed value, but instead follows a probability distribution. As described above, although some

delays may be acceptable for some applications (such as file transfers), unpredictable delays may pose a serious problem for real-time applications (such as streaming audio, video or voice communication applications). Consequently, these type applications can only run if the prior art systems have such a light traffic load that access to the channels is almost always immediate. It is desirable to design powerline networking systems such that real-time applications are supported even during high traffic volumes.

The present inventive MAC protocol method and apparatus solves these problems by providing "virtual circuits" between devices having real-time application requirements. In accordance with the present invention, a connection can be established between two devices wherein the connection guarantees a certain throughput between two points in the network. The virtual circuit connection also is guaranteed to have a constant average bit-rate and constant delay value. Therefore, the effect of the virtual circuit connection is that it is analogous to a hard-wired circuit connection, wherein the circuit connection has a throughput that is a fraction of that of the entire powerline network.

In accordance with the present MAC protocol method and apparatus, a virtual circuit can be created by establishing a periodic time slot that is reserved for use only by a specific selected transmitter. All of the other devices in the powerline network system must be aware of the reservation and must avoid making transmissions during this reserved time period. The selected transmitter can buffer its data during the time period leading up to the reserved time slot. For example, if the selected transmitter is providing a streaming video service, video information can be temporarily buffered during the time period leading up to the reserved time. The buffered data can then be transmitted in each reserved time period. Consequently, using the reservation-based access scheme of the present inventive MAC protocol method and apparatus, transmissions from the real-time oriented applications encounter a minimum possible delay. Moreover, the delays that are encountered using this approach advantageously remain constant over time.

Controller-less Reservation Based Access Method and Apparatus – An Exemplary Embodiment

An exemplary embodiment of the procedures that are used for controller-less reservation based access is described in detail in section 2.4 of the attached MAC specification, at pages 11-15. Those skilled in the art shall appreciate that alternatives to the embodiment described therein can be used without departing from the scope of the present invention. As described in

the MAC specification, the powerline networking system supports controller-less reservation based access modes to allow for the creation of virtual circuit connections that provide periodic, low latency, constant bit-rate service between an originating client and a destination client.

5

There are three MAC management related procedures related to controller-less reservation based access: (1) establishment of the reservation, (2) renewal of the reservation, and (3) termination of the reservation. In one embodiment of the present invention, the maximum reservation duration that can be established comprises either 256 periods, or 5 seconds, whichever is smaller. At the end of this time period, the reservation must either be renewed or terminated. The renewal process provides the MAC with the ability to change the payload format in response to changing physical medium conditions.

10

15

The management packets used to establish, renew, or terminate a reservation are transmitted as broadcast packets in ROBO mode. Because certain nodes may be able to receive from only one of the two clients involved in the reservation, the reservation information is transmitted by both clients that are party to the reservation.

Reservation Establishment

20

25

A reservation is established in the exemplary embodiment using a handshake process according to which the originating client establishes the reservation and the intended recipient acknowledges the establishment. The originating client initiates the process of establishing a reservation by broadcasting a ROBO mode packet that contains a reservation establishment (RE) payload. This RE payload informs all of the other clients (or devices) in the network of the time at which the reservation is to begin, the duration of the packets to be transmitted in the reservation, the period of transmission, and the lifetime of the reservation (that is, the number of packets that will be transmitted during the course of the reservation). The RE also provides a capability for establishing a reservation for a two-way circuit connection by allowing specification of a duration for a return transmission.

30

A two-way reservation comprises a forward transmission and a reverse transmission. In contrast, a one-way reservation comprises a forward transmission only. The forward transmission always occurs first in the embodiment described, and is transmitted by the originating client. If a reverse transmission exists, it occurs immediately after the forward

transmission completes, and is transmitted by the destination client. The reverse transmission must have the same period as the forward transmission, but it is not required to have the same duration as the forward transmission. In the embodiment described in the attached MAC specification, the maximum payload length that may be reserved for either the forward or the reverse transmission is 175 OFDM information symbols.

The reservation includes two time slots that are used to practice the present method and apparatus. A first time slot is provided in which a destination client can acknowledge the reservation. A second later time slot is provided in which clients can exchange broadcast messages that can either terminate or renew the reservation. In one embodiment, the time slot reserved for the initial reservation acknowledgment always begins 5 milliseconds after the start of the first OFDM symbol of the preamble of the packet containing the RE payload. The destination client broadcasts its reservation acknowledgment (RA) payload in a ROBO mode packet, repeating the fields describing the timing of the reservation for the benefit of clients that may have failed to receive the original reservation request.

If the originating client fails to receive the RA payload (which may happen due to a collision with its RE transmission or due to a collision with the RA), it assumes that the reservation has not been established and begins a new establishment procedure. The timing of the reservation establishment procedure is shown in the timing diagram of FIGURE 2.

As shown in FIGURE 2, once the RE and RA payloads have been exchanged, reservation access packets are exchanged according to pre-determined payload formats. In the embodiment described in the attached MAC specification, the receiving client does not acknowledge receipt of the reservation access packets.

As described above, other clients must not transmit during the reservation time. When a *MAC-data.req* for contention transmission is received by the MAC protocol layer, and the physical medium is determined to be available, the inventive MAC protocol layer must ensure that the physical medium remains available for the duration of the packet to be transmitted. If this is not the case, the MAC layer must enter a DEFER state and proceed as described in section 2.1.3 of the MAC specification.

Timing of the reserved slots is differential in nature. That is, each client predicts (or

computes) the start of the next reserved time based upon the end of the previously reserved time. When multiple reservations are simultaneously active, the client that established the first reservation serves as a reference for all subsequent reservation timing. Thus, the next reserved time for each active reservation is computed relative to the most recent transmission of the first reservation. When the first reservation concludes, the client that established the next reservation becomes the reference for further calculations. Clients that require different reservation periods from that being used by existing reservations must select a period such that no multiples of that period overlap reserved slots.

10 *Reservation Renewal and Termination*

The present MAC protocol method and apparatus allocates a time segment immediately following the last reservation access period during which reservations may be renewed or terminated. During this time segment, the two parties involved in a reservation use a three-way handshake to renew or terminate the reservation. This process (for bi-directional reservations) is shown in the message flow diagram of FIGURE 3a. FIGURE 3b shows a timing diagram used by the reservation renewal packets in accordance with the present invention. As shown in FIGURE 3b, if any reverse reservation access packets exist, a last reservation access packet 200 is transmitted at the time indicated. In the embodiment shown in FIGURE 3b and described in the attached MAC specification, a 63-micro-second interval 202 immediately following the last reservation access packet 200 of the reservation is reserved for the client that originated the reservation. The client that originated the reservation can use this interval 202 to begin transmission of a ROBO mode broadcast packet containing a reservation renewal (RR) payload message. No other client may attempt to transmit during this 63 microsecond period 202.

If the reservation is no longer needed, the RR payload terminates the reservation by setting the reservation lifetime field in the RR to zero. Otherwise it renews the reservation by providing new timing parameters for the reservation. If the reservation is bi-directional and the originating client determines that a different payload format should be used for reverse transmissions, it forms the packet with the RR payload to include a "PLLC" payload having new parameters to be used by the destination client. The timing information for the reservation must reflect any change in the length of the reservation access packet that will result from the new payload format.

When the destination client receives the RR payload, it responds with another broadcast RR payload. The time for this transmission is reserved so that no other client may transmit during this time and begins 500 microseconds after the conclusion of the last reservation access packet in the reservation. The destination client sets the reservation lifetime field of its RR payload to agree with that in the RR payload received from the originating client. If the reservation is being terminated (in one embodiment, indicated by the reservation lifetime being set to zero) then the handshake is completed once the destination client transmits its RR payload and there is no response from the originating client. If the reservation is being renewed and the destination client determines that a different payload format should be used for forward transmissions, then it includes a "PLLC" payload with the new parameters to be used by the originating client in the packet with the RR payload. The destination client updates the reservation timing information to reflect the time required to support reservation access packets that are formatted with the new payload format.

If the reservation is being renewed, the originating client transmits a broadcast RR payload when the originating client receives the RR from the destination client. The broadcast RR payload contains timing information that the originating client received from the destination client. The time that this transmission occurs is also reserved. That is, no other client may transmit during this time. In one exemplary embodiment, the transmission begins 1000 microseconds after the conclusion of the last reservation access packet in the reservation. This completes the handshake, and the reservation access proceeds as before. Finally, if a reservation is active in a network and a client (other than the originating and destination clients) fails to receive an RR that terminates or renews the reservation, that client must wait for three more cycles of the reservation before determining that the reservation is inactive.

Docket # 2009-0606-01

Error Conditions

In some cases, error conditions may make it impossible for a reservation to be established. Either client may reject the information contained in either an RE or an RR payload by setting the status field of the response payload to the appropriate value. In the embodiment described in the attached MAC specification, defined values for this field are:

- Status = 0: no error condition
- Status = 1: the reservation is rejected as specified because it interferes with other reservations known to the client sending the status
- Status = 2: the reservation is rejected because the receiver is not ready to accept the data
- Status = 3: a reservation length field exceeds 175 OFDM symbols
- Status = 4: the reservation is rejected for unspecified reasons.

Those skilled in the art shall recognize that alternative values for this field may be used without departing from the spirit or scope of the present invention.

Accordingly, the inventive MAC protocol method and apparatus provides a reservation based access technique in a powerline networking system having no central control mechanism. As described above with reference to FIGURES 2, 3a and 3b, an individual client may establish a reservation. One important aspect of this inventive technique is that all potentially interfering clients are prevented from interfering during the reservation establishment process. In powerline network systems, this task is complicated by the fact that not all clients can receive information from each other due to interference in the power lines.

In the scheme described above, each client maintains a list of active reservations. Timing for the reservations is differential and the master timing is derived from the transmissions of a client pair that first establish the reservation. As a consequence, each new transmission of a client re-establishes the timing reference.

As described above, all reservations are periodic. The period with which the reservation occurs and the amount of time reserved for each period are specified as part of the reservation establishment process. A first client wishing to establish a reservation with a second client must select a period, duration, and starting time for the reservation that are consistent with any existing reservations. The first client transmits a reservation establishment (RE) message to initiate the reservation. The RE contains the period, duration, and start time of the reservation.

The RE is transmitted in a broadcast mode with the intent that all clients should be able to receive it. These clients will then be aware of the timing of the reservation and will avoid transmitting during those times.

5 Due to the nature of the power line medium, there may be one or more other clients (a third client) that cannot receive transmissions from the first client but are in a position to interfere with the reception of the first client's messages by a second client. To inform these clients of the reservation timing, the second client thus transmits the reservation acknowledgment (RA) message. This message also contains the reservation timing information. The third client
10 should be able to receive this transmission and thereby will be informed of the reservation timing.

As described above, reservations have a maximum allowed lifetime, and after this time they must either be terminated or renewed. This is an important aspect of the present inventive
15 technique. Without maximum allowed lifetimes it would be possible for clients to interpret reservations as having indefinite (perhaps infinite) lifetimes if the clients miss a control message ending the reservation. As described above, the reservation is renewed using RR messages. The RR message operates as a three-way handshake that allows it to be used to renegotiate physical layer parameters based on changing channel conditions. These physical
20 layer parameters might include the modulation type, the error correction coding, and for OFDM systems, the set of frequencies that are used for transmission.

If a client that is not part of the reservation fails to receive the RR messages that either terminate or extend the reservation, it must assume that the reservation has been extended to
25 ensure that there is no collision. Of course, if the reservation was in fact terminated, this client also will not receive an RR after the next lifetime. Per the rules of the protocol, if a client misses three consecutive RR opportunities, it assumes that the reservation has terminated and that the terminating RR was missed. The client then can use the channel during that reserved time.

30 As described above, another important aspect of the inventive reservation technique is the ability to establish bi-directional reservations. This is a very useful feature for telephony access. Another important aspect of the reservation scheme is that only the initial RE is transmitted in a contention access mode. In other words, when the RE is transmitted, no

reservation has been established and the RE consequently may collide with some other client transmissions and therefore may not be received. However, once the RE has been transmitted without collision, a time slot is implicitly provided for the RA and also for the subsequent RR transmissions so that these messages do not have to contend. This makes the messages far more reliable and also makes their use of the channel more efficient because there is no need to re-transmit these messages due to collisions.

A last important aspect of the present inventive reservation scheme is that the other clients can segment their contention-based access messages to conform to the reservations. In other words, if a third client queues a 1 millisecond duration message, for transmission 500 microseconds before the start of a reserved time, the inventive protocol allows the client to divide the message into smaller segments. This allows the third client to transmit part of the message in the 500 microseconds before the reservation and to transmit the rest of the message after the reservation ends. The receiver can then reassemble the original message. This aspect of the reservation scheme improves the overall efficiency of the reservation based access scheme of the present invention.

Logical Network Identifiers (LNI)

As noted above in the background of the invention, the present MAC protocol method and apparatus includes means for uniquely identifying logical networks in the powerline networking system. In accordance with the present method and apparatus, devices (also referred to herein as clients) may only be members of a single logical network. That is, although clients may be electrically (and physically) coupled to the same physical medium as are other clients, they can be treated differently by the present MAC protocol invention using the concept of logical networks. A client is said to belong to one, and only one, logical network. As a consequence, a client may only exchange data with other member clients belonging to its logical network. The present inventive MAC protocol uses logical network identifiers (LNI) to uniquely identify the logical networks in the system.

The clients can use the LNI information to determine if they should attempt to receive a given packet. One approach to transmitting the LNI information is to include the LNI in each packet that is transmitted. Alternatively, the LNI can be communicated using a management message in which the transmitting client can declare which LNI it belongs to. Other clients can then build a table that maps client addresses to LNIs. If the format of typical messages includes the

address of the source client (as is the case here), then the source address can be used to determine which LNI the client belongs to.

5 In an ideal world, each logical network sharing a physical medium would have a unique identifier. The number of bits required to represent this identifier is a function of the number of logical networks that can share a physical medium. In one proposed powerline networking system, this number is 128. This would mean that at least 7 bits would be required to represent the LNI.

10 In systems such as powerline networking systems, where no coordinating management exists, there is no convenient way to ensure that LNIs are not accidentally re-used. As described above, it is not practical (or desirable) for a first user of devices in a first house to ask a second user in a second neighboring house which LNI the second user might have selected, and *vice versa*. One means for addressing this problem is to ensure that there are far more possible
15 LNIs than the maximum number of logical networks, and then to select the LNIs in a random manner. This approach is similar to the approach used for garage door openers or house keys: there is a finite number of possible keys, but the number is sufficiently large that it is very unlikely that two parties will select the same key settings.

20 The problem then is to determine which method to use to assign the random values. The present invention solves this problem by using a "password-like" value that is entered by the logical network owner. The inventive method then uses a hash function to map this value to the LNI. In one exemplary embodiment described herein, the street address of a network owner is used for the password value. The network owner would enter this street address
25 information into the powerline networking system during installment of the system. The characters in this address are then converted into a bit sequence using a well known ASCII mapping technique. This bit sequence is then used to create a 32-bit Cyclic Redundancy Check (CRC) (alternative lengths can be selected) which then serves as the LNI. One exemplary embodiment of the LNI of the present invention is described in more detail in
30 section 2.3.1, at pages 10-11, of the attached MAC specification. This exemplary embodiment is now described.

As described in the attached MAC specification, the logical network is identified using a 32-bit LNI field in the beacon payload. A network name of any length may be used by the

management entity. The network name is compressed to a 32-bit LNI using a 32-bit CRC generator 300 as shown in FIGURE 4. The CRC polynomial represented by FIGURE 4 is given by Equation 1 below:

5 **Equation 1:**

$$g(x) = x^{32} + x^{26} + x^{23} + x^{22} + x^{16} + x^{12} + x^{11} + x^{10} + x^8 + x^7 + x^5 + x^4 + x^2 + x^1 + 1$$

In the embodiment described in the attached MAC specification, the LNI is formed as follows.

Registers in the CRC are first initialized to zero. The two switches are set to an "UP" position, and the ASCII translated bits of the network name are input (one at a time) into the CRC generator 300 of FIGURE 4. When all of the bits of the network name have been input to the CRC generator 300, the two switches are moved to the "DOWN" position and the CRC generator is clocked 32 times, producing an output bit at each clock. The first output bit is the least significant bit (LSB) of the LNI. The last output bit is the most significant bit (MSB) of the LNI.

With a 32-bit CRC, the probability that two addresses will map to the same 32-bit LNI value is $\frac{1}{2^{32}}$, or about one in four billion. If there are 128 logical networks sharing the same physical medium (as has been proposed in one powerline networking system), then the probability that they will not all choose different LNI values is given by the solution to the well known "birthday problem" given below in Equation 2:

Equation 2: $p(m \text{ unique lni values with } N \text{ total possible}) = 1 - \frac{N!}{N^m \cdot (N - m)!}$

For a 32-bit CRC and $m=128$, this yields a value of about 2 in one million. This probability is sufficiently small to be satisfactory in most system configurations, especially given how infrequently the physical medium will actually be required to support the maximum number of logical networks allowed. Thus, an inventive method of assigning unique LNIs has been described. The inventive method facilitates generation of LNIs with very low probability of non-uniqueness. The method assigns the unique LNIs in a system that has no central controlling authority and that has multiple logical networks sharing a physical medium.

Client devices that lack user input/output (I/O) capability must obtain the LNI from some other device that is connected to the network. This occurs in the process of encryption key distribution described below. In the embodiment described in the attached MAC specification, an "all-zero" LNI is reserved for use by clients that must receive the LNI from another client in this manner. The inventive method of distributing encryption keys, especially for devices that have no user I/O capability, is now described.

Encryption Key Distribution

Encryption systems typically comprise two major components: an *algorithm* used to operate on data and a *key* system used to initialize the algorithm. The key system initializes the algorithm so that both the transmitter and the receiver use the algorithm in a manner that allows received data to be deciphered. Most encryption systems rely on the encryption keys for security. Most systems assume that the encryption algorithm is known (or can be derived) by a potential attacker. Prior art cryptography techniques are described in a text by Bruce Schneier, entitled "*Applied Cryptography*", published by Wiley and Sons in 1996, and hereby incorporated by reference herein for its teachings on cryptography.

Application of cryptography techniques to data networks has primarily been a problem of designing appropriate methods for the secure management of encryption keys. There are relatively few well-tested encryption algorithms currently being used. Most cryptography designs select one of these well-known algorithms. The requirements of the selected key management algorithm depend heavily on the topology of the network. As described above, in the powerline networking system environment, there typically are multiple logical networks sharing a single physical medium. As described above, logical networks are defined as having a group of member clients who are intended to share data. Each member of a logical network therefore "trusts" the other members of the logical network. That is, the assumption is that no member of the logical network will improperly use data that is accessible from other members of the logical network. In practical terms, in powerline networking systems, a logical network might comprise all of the network devices connected together in a given home.

As described above, one of the drawbacks associated with powerline networking systems is that the physical medium (*i.e.*, the power lines) is typically shared between multiple households. Indeed, in a typical configuration, power lines are shared between 5-8 homes in single family dwelling areas, and between more homes in apartment dwellings. Consequently,

the multiple homes are able to receive each other's physical layer transmissions. In the absence of some type of encryption technique, each household's data is vulnerable to access by potential attackers in a neighboring house or apartment. In addition, because homes often have exterior outlets as described above, data is also vulnerable to attack by potential intruders via access using the exterior outlets.

Thus an encryption system is needed in powerline networking systems in which data can be safely shared among members of a given logical network and is protected from parties that are not members of the logical network. One straightforward means for managing encryption keys in powerline networking systems is to have the keys manually entered into each client of the logical network. If each client has the same encryption key, the key can simply be derived from a logical network password, which can then be converted into a fixed length encryption key using one of several well-known techniques. For example, a hash function can be used for this purpose wherein the hash function accepts variable length inputs and convert the inputs into fixed length outputs via some difficult-to-reverse algorithm. This approach works quite well, however its applicability to powerline networking systems is limited in situations where some or all of the devices lack a means by which the encryption key can be manually entered. Exemplary devices lacking user I/O capability are network printers and gateway devices (such as cable modems) that allow multiple PCs to share a broadband access device. The present inventive encryption key management method and apparatus provides an encryption key management solution in powerline networking systems, wherein some or all of the devices have limited, or no, user I/O capability.

One exemplary embodiment of the present inventive encryption key management method and apparatus is described in detail in the attached MAC specification in section 2.7, at pages 16-18. The details of the format of the MAC management message used for key update (referred to in the attached MAC specification as the "Encryption key update payload") are given in section 3.1.2.1.1.8, at pages 30-31. This exemplary embodiment is now described. However, those skilled in the encryption key design art shall appreciate that alternative encryption key management embodiments can be used without departing from the spirit or scope of the present invention.

Procedures for Encryption – An Exemplary Embodiment

5 The powerline networking system security protocol essentially serves three system goals. First, the security protocols used help to ensure the confidentiality of the network. Data transmitted on the physical medium is accessible only to authorized entities. All members of a common logical network are regarded as authorized entities. Second, the security protocol provides for secure key management. The security of encryption keys is maintained. Finally, the security protocol provides an ability to upgrade the encryption algorithms used.

10 In the embodiment described in the attached MAC specification, the security protocol is intended to ensure that data is known only by the source and destination clients. However, the security protocol does not provide a "non-repudiation" function. That is, receipt of a message from a client does not irrevocably prove that it came from the apparent sender. The security protocol also does not provide for protection against the monitoring of the volume of data that is exchanged in the network. Nor does the protocol protect against attacks that
15 disrupt the network through the use of spurious control messages.

Encryption algorithm

20 The baseline encryption algorithm for the powerline networking system of the exemplary embodiment is a Data Encryption Standard (DES) operating as a stream cipher in output feedback mode. One skilled in the encryption art shall recognize that the other algorithms and modes can be utilized with the key management algorithm of the present invention without departing from the scope or spirit of the present invention. The keystream is applied only to the payload bits as indicated in section 3 of the attached MAC specification.

25 In accordance with the present invention, each client indicates in its beacon payload message which encryption algorithms it supports. A receiving client can indicate which algorithm it desires the source client to use using a "PLLC" payload message as defined in the attached MAC specification. However, the receiving client must select an algorithm that is supported by the source client.

30 Procedures for Encryption Key Management

As described in the attached specification, all clients that are members of the same logical network must use the same encryption key. It is generally assumed that key changes will be made infrequently in the proposed powerline networking systems. The key is typically

provided by a client host to the MAC layer using a "MAC-KEY.req" primitive. As described above, in some cases, some or all of the client devices may lack a user I/O interface suitable for manually entering encryption keys. In these cases, the present invention provides a means for allowing these non-I/O capability clients to use a key received from another client via the physical medium.

A client device lacking user I/O capability must have a key (e.g., default key and "hard-programmed" key) upon entering the network. This key is independent of the network that is associated with the client device. A system user installing the client device can know either the key or the password utilized to create the key. The key typically differs from the key used by other members of the network, and is used only to enable the device to receive the key being used by the other members of the network. In the exemplary embodiment the client always retains this key.

The client device that requires the network key (the "receiving client") obtains the key by receiving an encryption key update payload message from any other device (the "originating client") in the network. One exemplary embodiment of the encryption key update payload message is described in section 3.1.2.1.1.8, at pages 30-31, of the attached MAC specification. As described therein, this payload contains the encryption key currently being used by the network, as encrypted using the hard-programmed key of the receiving client and using an initialization vector contained in the key update payload. In order to minimize the possibility that a Forward Error Correction (FEC) decoder error may erroneously produce a key update, a 32-bit CRC is included in the key update payload. Using the CRC generator 300 of FIGURE 4, the CRC is formed over the entire payload prior to encryption. The encryption is applied to the key field and to the CRC field.

The receiving client acknowledges receipt of the encryption key by transmitting a key update acknowledgment payload message. To create this payload message, the receiving client selects a new initialization vector, fills the encryption key field with its hard-programmed key value, and computes a new CRC over the entire payload. It then encrypts the key field and the CRC using information contained in an initialization vector (IV) that is returned in the key update acknowledgment and the key that it previously obtained from the key update message. Henceforth all encrypted fields are encrypted using the key that was received in the key update message.

5 If the originating client does not receive the key update acknowledgment, it must re-transmit the key update after each beacon it receives from the receiving client. If the receiving client has a network key, but receives another key update message, it must replace the network key currently in use with the network key contained in the key update message. It then must acknowledge the key update message as described above.

10 In the embodiment described in the attached MAC protocol specification, the key update payload also contains the LNI currently being used by the logical network. The receiving client accepts the LNI only if the CRC passes. The procedures used by the present invention for encryption synchronization are described in detail in section 2.7.3, at pages 17 and 18 of the attached MAC specification.

15 Thus, an inventive technique for generating and managing encryption keys in powerline networking systems has been described. In practical terms, the above-described technique can be implemented as follows. A device that lacks user I/O capability can be provided with a hard-wired encryption key or default key, which can be set to any possible value. This hard-wired key might be printed on a label that is applied to the packaging of an installation CDR (or other readable medium) that is shipped with the device. When a user installs the device,
20 the user can load the CDR into a networked PC or other device that does have user I/O capability (e.g., a DVD player having a TV and remote control). An installation program on the CDR can be programmed to ask the user for the hard-wired key (wherein the key or the password that gets hashed to create the key is printed on a sticker accompanying the packaging of the device being installed). The installation program can also ask the user for
25 the logical network password. The inventive MAC protocol layer of the powerline network protocol then uses the sticker key to encrypt the current logical network key as described above. The inventive MAC protocol can then transmit this encrypted key to the new device. The logical network key is thus securely passed to the new device, and all other members of the logical network can now exchange encrypted data with the device. If the device loses the
30 logical network key, or if the key changes, another device can re-transmit the key using the same MAC management message originally used to provide the key.

In some ways, it might appear that a more straightforward approach may be to simply use the hard-programmed key for all of the exchanges with the new device. Disadvantageously, this

approach has some shortcomings. The difficulty with this approach is that all of the other devices in the logical network must be aware of this key as well. Consequently, the key must either be manually loaded into each device (which may not be possible if the devices do not all have user I/O capability), or the key must be transmitted to the devices using the network.

5 This latter approach is more complicated than simply performing a single key exchange with the new device. In addition, using this approach, all of the other devices are required to maintain separate key information for each device in the logical network. Therefore, although at first glance more straightforward, it is probably more complicated to simply use the hard-programmed key for all of the exchanges with the new device.

10

In summary, the inventive encryption key management process comprises four main steps. The main steps can be summarized as follows:

- 15 1. A first device that does not have user I/O capability has a hard-programmed key assigned to it. The hard-programmed key is used only for an initial key exchange. The first device comprises non-volatile storage or similar memory means for purposes of holding the logical network key.
- 20 2. Upon addition of the first device into the logical network, the hard-programmed key is entered into another device member of the logical network (a second device) that does have user I/O capability.
- 25 3. The second device then transmits a MAC management message for key update to the first client. This MAC management message contains the currently used logical network key as encrypted using the hard-programmed key of the first device.
4. Upon receipt of the MAC management message, the first device loads the logical network key into a non-volatile storage means. The first device then uses this encryption key for all subsequent encryption operations in the logical network, with the exception that if it receives another key update MAC management message, it will then use its hard-programmed key to decipher.

30 Although an exemplary embodiment of the encryption key assignment and management method of the present invention has been described, those skilled in the art shall recognize that modifications can be made to the described embodiment without departing from the spirit and scope of the present invention. For example, an alternative means can be used to store the logical network key in the first device without departing from the scope of the present invention. Instead of using a non-volatile storage means, an alternative memory means could be used to store the key information. Further, in an alternative embodiment, the key can be generated without use of a CRC. Alternatively, alternative size CRCs can be used to generate the key. These and other modifications can be made to the exemplary embodiment without departing from the scope of the present invention.

40

Summary

In summary, an inventive Medium Access Control protocol for powerline networking systems has been described. The inventive MAC protocol controls access to and use of a physical medium (power lines) in a powerline networking system. In one embodiment, the MAC protocol layer of the powerline networking system uses a Carrier Sense Multiple Access (CSMA) protocol with modifications to support special requirements for applications requiring low latency. The protocol supports both contention-based access and reservation-based access. Reservation-based access can operate in either a controller-less mode or in a mode wherein a network controller is present.

A complete MAC protocol is described that specifically addresses concerns unique to home powerline networking systems. In one embodiment, the protocol is intended to operate with a physical layer that uses an OFDM modulation scheme. However, the inventive MAC protocol method and apparatus is contemplated for use with physical layers using other types of modulation schemes. The inventive MAC protocol method and apparatus includes a method of providing "blanking intervals" in which devices using newer versions of the protocol can "clear out" earlier version devices. The use of blanking intervals greatly eases backward compatibility when the protocol is upgraded to newer versions. The method of using blanking intervals is closely coupled to another inventive technique of using "beacons" that propagate the blanking timing information throughout the network. Using the inventive beacon method and apparatus, devices are informed as to whether the blanking information has expired. The inventive MAC method and apparatus also includes a method of establishing and maintaining "virtual circuit" connections between selected clients. Virtual circuits can be established in the powerline networking system even when the networking system does not have a central controller.

The inventive MAC protocol method and apparatus also provides a facility for assigning unique Logical Network Identifiers (LNIs) to logical networks in the powerline networking system. The LNIs uniquely identify each of the logical networks in the network. The LNIs are generated even in systems where no central control mechanism is used. Finally, the inventive MAC protocol method and apparatus includes a means for creating, managing, and distributing network encryption keys. The encryption keys are used by the devices in the powerline networking system to prevent data from being shared with unauthorized users. A

method for distributing encryption keys to devices not having user input/output capability is described.

- 5 A number of embodiments of the present invention have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of the invention. Accordingly, it is to be understood that the invention is not to be limited by the specific illustrated embodiment, but only by the scope of the appended claims.

090909-12192060